UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
10/603,881	06/25/2003	David Darden Chambliss	IBMS.004PA(0523)	8733

62627 7590 03/06/2007

DAVID W. LYNCH
CHAMBLISS, BAHNER & STOPHEL
1000 TALLAN SQUARE-S
TWO UNION SQUARE
CHATTANOOGA, TN 37402

EXAMINER	
WON, MICHAEL YOUNG	
ART UNIT	**PAPER NUMBER**
2155	

SHORTENED STATUTORY PERIOD OF RESPONSE	MAIL DATE	DELIVERY MODE
3 MONTHS	03/06/2007	PAPER

Please find below and/or attached an Office communication concerning this application or proceeding.

If NO period for reply is specified above, the maximum statutory period will apply and will expire 6 MONTHS from the mailing date of this communication.

Office Action Summary	Application No.	Applicant(s)	
	10/603,881	CHAMBLISS ET AL.	
	Examiner	Art Unit	
	Michael Y. Won	2155	

-- The MAILING DATE of this communication appears on the cover sheet with the correspondence address --

Period for Reply

A SHORTENED STATUTORY PERIOD FOR REPLY IS SET TO EXPIRE 3 MONTH(S) OR THIRTY (30) DAYS, WHICHEVER IS LONGER, FROM THE MAILING DATE OF THIS COMMUNICATION.

- Extensions of time may be available under the provisions of 37 CFR 1.136(a). In no event, however, may a reply be timely filed after SIX (6) MONTHS from the mailing date of this communication.
- If NO period for reply is specified above, the maximum statutory period will apply and will expire SIX (6) MONTHS from the mailing date of this communication.
- Failure to reply within the set or extended period for reply will, by statute, cause the application to become ABANDONED (35 U.S.C. § 133). Any reply received by the Office later than three months after the mailing date of this communication, even if timely filed, may reduce any earned patent term adjustment. See 37 CFR 1.704(b).

Status

1)☒ Responsive to communication(s) filed on *25 June 2003*.

2a)☐ This action is **FINAL**. 2b)☒ This action is non-final.

3)☐ Since this application is in condition for allowance except for formal matters, prosecution as to the merits is closed in accordance with the practice under *Ex parte Quayle*, 1935 C.D. 11, 453 O.G. 213.

Disposition of Claims

4)☒ Claim(s) *1-24* is/are pending in the application.

4a) Of the above claim(s) _____ is/are withdrawn from consideration.

5)☐ Claim(s) _____ is/are allowed.

6)☒ Claim(s) *1-3,5,7-9,11,13-15,17,18 and 20-24* is/are rejected.

7)☐ Claim(s) *4,6,10,12,16,19 and 21* is/are objected to.

8)☐ Claim(s) _____ are subject to restriction and/or election requirement.

Application Papers

9)☐ The specification is objected to by the Examiner.

10)☐ The drawing(s) filed on _____ is/are: a)☐ accepted or b)☐ objected to by the Examiner.

Applicant may not request that any objection to the drawing(s) be held in abeyance. See 37 CFR 1.85(a).

Replacement drawing sheet(s) including the correction is required if the drawing(s) is objected to. See 37 CFR 1.121(d).

11)☐ The oath or declaration is objected to by the Examiner. Note the attached Office Action or form PTO-152.

Priority under 35 U.S.C. § 119

12)☐ Acknowledgment is made of a claim for foreign priority under 35 U.S.C. § 119(a)-(d) or (f).

a)☐ All b)☐ Some * c)☐ None of:

1.☐ Certified copies of the priority documents have been received.

2.☐ Certified copies of the priority documents have been received in Application No. _____.

3.☐ Copies of the certified copies of the priority documents have been received in this National Stage application from the International Bureau (PCT Rule 17.2(a)).

* See the attached detailed Office action for a list of the certified copies not received.

Attachment(s)

1) ☒ Notice of References Cited (PTO-892)
2) ☐ Notice of Draftsperson's Patent Drawing Review (PTO-948)
3) ☒ Information Disclosure Statement(s) (PTO/SB/08) Paper No(s)/Mail Date *6/25/03*.
4) ☐ Interview Summary (PTO-413) Paper No(s)/Mail Date. _____ .
5) ☐ Notice of Informal Patent Application
6) ☐ Other: _____.

DETAILED ACTION

1. This action is in response to the application filed June 25, 2003.

2. Claims 1-24 have been examined and are pending with this action.

Claim Objections

3. Claim 21 is objected to because of the following informalities: There is a typographical error in page 35, line 6 ("lest" should be replaced by "least"). Appropriate correction is required.

Claim Rejections - 35 USC § 112

The following is a quotation of the second paragraph of 35 U.S.C. 112:

> The specification shall conclude with one or more claims particularly pointing out and distinctly claiming the subject matter, which the applicant regards as his invention.

4. Claims 8, 18, 23 are rejected under 35 U.S.C. 112, second paragraph, as being indefinite for failing to particularly point out and distinctly claim the subject matter which applicant regards as the invention.

Claim 8 recites the limitation "the at least one I/O performance gateway" in pg.31, lines 8-9. There is insufficient antecedent basis for this limitation in the claim.

Claim 18 recites the limitation "the SLA service module" in pg.34, line 6. There is insufficient antecedent basis for this limitation in the claim.

Claim 21 recites the limitation " the SLA service module " in pg.35, line 4. There is insufficient antecedent basis for this limitation in the claim.

Claim 23 recites the limitations "the means for storing" in pg.38, line 2 and "the service level agreement and policy" in pg.39, lines 4-5. There is insufficient antecedent basis for this limitation in the claim.

Claim Rejections - 35 USC § 103

The following is a quotation of 35 U.S.C. 103(a) which forms the basis for all obviousness rejections set forth in this Office action:

> (a) A patent may not be obtained though the invention is not identically disclosed or described as set forth in section 102 of this title, if the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art to which said subject matter pertains. Patentability shall not be negatived by the manner in which the invention was made.

5. Claims 1-3, 5, 7-9, 11, 13-15, 17, 18 and 20-24 are rejected under 35 U.S.C. 103(a) as being unpatentable over Guha (US 2002/0194324 A1) in view of DeKoning et al. (US 6,457,098 B1).

INDEPENDENT:

As per **claim 1**, Guha teaches a storage area network, comprising:

a storage system for providing storage of system data (see Fig.5, #58; pg.2, [0013]: "storage center"; and pg.5, [0065]: "content storage 58");

at least one application host running an application thereon and accessing data on the storage system (see Fig.5, #62; and pg.5, [0064]: "preferably contains application servers 62");

at least one I/O performance gateway, disposed between the at least one application host and the storage subsystem, for intercepting I/O operations (see Fig.4, #42; and pg.5, [0063]: "The plurality of servers communicate with storage device 44 through a network storage switch 42 (i.e., SAN switch)");

a database for storing service level agreements (see pg.3, [0039]: "SLA Enforcement device"); and

a Service Level Agreement (SLA) server for controlling the at least one I/O performance gateways based on the service level agreements (see pg.5, [0068]: "Based on the policy specified in the QoS Enforcer 34, content requests may be... admitted into data center or rerouted to another data center if its SLA cannot be net at the current site"), wherein the at least one I/O performance gateway sending statistic data to the SLA server (see pg.4, [0049]: "The resource status preferably contains information received from the content controller 36 which indicates the status of the components which retrieved the data") and taking requests from the SLA server to control I/O operations (pg.5, [0069]: "communicated by the QoS Enforcer 34 to the content controller 36");

wherein the SLA server further comprises:

a database manager for maintaining connections to the database wherein service level agreements and performance data are maintained (see pg.3, [0068]: "Based on the policy specified in the QoS Enforcer 34, content request may be dropped or requeued");

an SLA services module for analyzing data and controlling actions based on the service level agreements and policy (see pg.4, [0045]: "The routing decisions are determined through a combination of preset QoS policy and current expected load at the data center");

an application server for communicating with clients to display monitoring information and for communication with the database manager (see Fig.5, #62); and

a performance monitor for communicating with the at least one I/O performance gateway to collect data and send throttling requests based upon signals from the SLA services module (see pg.3, [0068]: "Based on the policy specified in the QoS Enforcer 34, content request may be dropped or requeued").

Guha does not explicitly teach wherein the performance monitor is configured to provide at least one thread pair for processing inbound signals from the at least one I/O device/module being sent to the services module via an inbound thread and for processing outbound signals to the at least one I/O device/module received from the services module via an outbound thread, the inbound thread and the outbound thread operating asynchronously to provide non-blocking messaging.

DeKoning teaches a performance monitor is configured to provide at least one thread pair for processing inbound signals from the at least one I/O device/module being sent to the services module via an inbound thread and for processing outbound signals to the at least one I/O device/module received from the services module via an outbound thread, the inbound thread and the outbound thread operating

asynchronously to provide non-blocking messaging (see col.17, lines 8-16: "multi-threaded programming technique may be employed within a RAID controller to permit other processing to continue while element 700 is operable to await receipt of a new I/O request").

It would have been obvious to a person of ordinary skill in the art at the time the invention was made to modify the system of Guha in view of DeKoning so that performance monitor is configured to provide at least one thread pair for processing inbound signals from the at least one I/O device/module being sent to the services module via an inbound thread and for processing outbound signals to the at least one I/O device/module received from the services module via an outbound thread. One would be motivated to do so because Guha teaches of dynamically balancing input/output loads on the storage center (see pg.2, [0018]).

As per **claim 8**, Guha teaches a Service Level Agreement (SLA) server, comprising:

a database manager for maintaining connections to the database wherein service level agreements and performance data are maintained (see pg.3, [0068]: "Based on the policy specified in the QoS Enforcer 34, content request may be dropped or requeued");

an SLA services module for analyzing data and controlling actions based on the service level agreements and policy (see pg.4, [0045]: "The routing decisions are

determined through a combination of preset QoS policy and current expected load at the data center");

an application server for communicating with clients to display monitoring information and for communication with the database manager (see Fig.5, #62); and

a performance monitor for communicating with the at least one I/O performance gateway to collect data and send throttling requests based upon signals from the SLA services module (see pg.3, [0068]: "Based on the policy specified in the QoS Enforcer 34, content request may be dropped or requeued").

Guha does not explicitly teach wherein the performance monitor is configured to provide at least one thread pair for processing inbound signals from the at least one I/O device/module being sent to the services module via an inbound thread and for processing outbound signals to the at least one I/O device/module received from the services module via an outbound thread, the inbound thread and the outbound thread operating asynchronously to provide non-blocking messaging.

DeKoning teaches a performance monitor is configured to provide at least one thread pair for processing inbound signals from the at least one I/O device/module being sent to the services module via an inbound thread and for processing outbound signals to the at least one I/O device/module received from the services module via an outbound thread, the inbound thread and the outbound thread operating asynchronously to provide non-blocking messaging (see col.17, lines 8-16: "multi-threaded programming technique may be employed within a RAID controller to permit

other processing to continue while element 700 is operable to await receipt of a new I/O request").

It would have been obvious to a person of ordinary skill in the art at the time the invention was made to modify the system of Guha in view of DeKoning so that performance monitor is configured to provide at least one thread pair for processing inbound signals from the at least one I/O device/module being sent to the services module via an inbound thread and for processing outbound signals to the at least one I/O device/module received from the services module via an outbound thread. One would be motivated to do so because Guha teaches of dynamically balancing input/output loads on the storage center (see pg.2, [0018]).

As per **claim 14**, Guha teaches a performance monitor for controlling communication between two functional entities, comprising:

at least one processor configured for processing signals between at least one I/O performance gateway and SLA services module processes (pg.5, [0069]: "communicated by the QoS Enforcer 34 to the content controller 36").

Guha does not explicitly teach a thread pair associated with each of the at least one processors, each thread pair processing inbound signals from the at least one I/O device/module being sent to the services module via an inbound thread and for processing outbound signals to the at least one I/O device/module received from the services module via an outbound thread, the inbound thread and the outbound thread operating asynchronously to provide non-blocking messaging.

DeKoning teaches a thread pair associated with each of the at least one processors, each thread pair processing inbound signals from the at least one I/O device/module being sent to the services module via an inbound thread and for processing outbound signals to the at least one I/O device/module received from the services module via an outbound thread, the inbound thread and the outbound thread operating asynchronously to provide non-blocking messaging (see col.17, lines 8-16: "multi-threaded programming technique may be employed within a RAID controller to permit other processing to continue while element 700 is operable to await receipt of a new I/O request").

It would have been obvious to a person of ordinary skill in the art at the time the invention was made to modify the system of Guha in view of DeKoning by implementing a thread pair associated with each of the at least one processors, each thread pair processing inbound signals from the at least one I/O device/module being sent to the services module via an inbound thread and for processing outbound signals to the at least one I/O device/module received from the services module via an outbound thread. One would be motivated to do so because Guha teaches of dynamically balancing input/output loads on the storage center (see pg.2, [0018]).

As per **claim 18**, Guha teaches a program storage device readable by a computer, the medium tangibly embodying one or more programs of instructions executable by the computer to perform a method for providing non-blocking, minimum threaded two-way messaging, the method comprising:

providing at least one processor for controlling communication between SLA processes for the SLA service module and at least one I/O performance gateway (pg.5, [0069]: "communicated by the QoS Enforcer 34 to the content controller 36").

Guha does not explicitly teach providing a thread pair associated with each of the at least one processors for processing inbound signals from the at least one I/O device/module being sent to the services module via an inbound thread and for processing outbound signals to the at least one I/O device/module received from the services module via an outbound thread, the inbound thread and the outbound thread operating asynchronously to provide non-blocking messaging.

DeKoning teaches providing a thread pair associated with each of the at least one processors for processing inbound signals from the at least one I/O device/module being sent to the services module via an inbound thread and for processing outbound signals to the at least one I/O device/module received from the services module via an outbound thread, the inbound thread and the outbound thread operating asynchronously to provide non-blocking messaging (see col.17, lines 8-16: "multi-threaded programming technique may be employed within a RAID controller to permit other processing to continue while element 700 is operable to await receipt of a new I/O request").

It would have been obvious to a person of ordinary skill in the art at the time the invention was made to modify the system of Guha in view of DeKoning by implementing providing a thread pair associated with each of the at least one processors for processing inbound signals from the at least one I/O device/module being sent to the

services module via an inbound thread and for processing outbound signals to the at least one I/O device/module received from the services module via an outbound thread. One would be motivated to do so because Guha teaches of dynamically balancing input/output loads on the storage center (see pg.2, [0018]).

As per **claim 21**, Guha teaches a method for providing non-blocking, minimum threaded two-way messaging, comprising:

providing at least one processor for controlling communication between SLA processes for the SLA service module and at least one I/O performance gateway (pg.5, [0069]: "communicated by the QoS Enforcer 34 to the content controller 36").

Guha does not explicitly teach providing a thread pair associated with each of the at least one processors for processing inbound signals from the at least one I/O device/module being sent to the services module via an inbound thread and for processing outbound signals to the at least one I/O device/module received from the services module via an outbound thread, the inbound thread and the outbound thread operating asynchronously to provide non-blocking messaging.

DeKoning teaches providing a thread pair associated with each of the at least one processors for processing inbound signals from the at least one I/O device/module being sent to the services module via an inbound thread and for processing outbound signals to the at least one I/O device/module received from the services module via an outbound thread, the inbound thread and the outbound thread operating asynchronously to provide non-blocking messaging (see col.17, lines 8-16: "multi-

threaded programming technique may be employed within a RAID controller to permit other processing to continue while element 700 is operable to await receipt of a new I/O request").

It would have been obvious to a person of ordinary skill in the art at the time the invention was made to modify the system of Guha in view of DeKoning by implementing providing a thread pair associated with each of the at least one processors for processing inbound signals from the at least one I/O device/module being sent to the services module via an inbound thread and for processing outbound signals to the at least one I/O device/module received from the services module via an outbound thread. One would be motivated to do so because Guha teaches of dynamically balancing input/output loads on the storage center (see pg.2, [0018]).

As per **claim 22**, Guha teaches a storage area network comprising:

storage means for providing storage of system data (see Fig.5, #58; pg.2, [0013]: "storage center"; and pg.5, [0065]: "content storage 58");

at least one application means for running an application thereon and accessing data on the storage means (see Fig.5, #62; and pg.5, [0064]: "preferably contains application servers 62");

at least one gateway means, disposed between the at least one application host and the storage subsystem, for intercepting I/O operations (see Fig.4, #42; and pg.5, [0063]: "The plurality of servers communicate with storage device 44 through a network storage switch 42 (i.e., SAN switch)");

means for storing service level agreements (see pg.3, [0039]: "SLA Enforcement device"); and

means for controlling the at least one I/O performance gateways based on the service level agreements (see pg.5, [0068]: "Based on the policy specified in the QoS Enforcer 34, content requests may be... admitted into data center or rerouted to another data center if its SLA cannot be net at the current site"), wherein the at least one gateway means sends statistic data to the means for controlling (see pg.4, [0049]: "The resource status preferably contains information received from the content controller 36 which indicates the status of the components which retrieved the data") and takes requests from the means for controlling to control I/O operations (pg.5, [0069]: "communicated by the QoS Enforcer 34 to the content controller 36");

wherein the means for controlling further comprises:

means for managing and maintaining connections to the means for storing service level agreements (see pg.3, [0068]: "Based on the policy specified in the QoS Enforcer 34, content request may be dropped or requeued");

means for analyzing data and controlling actions based on the service level agreements and policy (see pg.4, [0045]: "The routing decisions are determined through a combination of preset QoS policy and current expected load at the data center");

mean for communicating with clients to display monitoring information and for communication with the means for managing and maintaining connections to the means for storing service level agreements (see Fig.5, #62); and

means for communicating with the at least one gateway means to collect data and send throttling requests based upon signals from the means for analyzing data and controlling actions (see pg.3, [0068]: "Based on the policy specified in the QoS Enforcer 34, content request may be dropped or requeued").

Guha does not explicitly teach wherein the means for communicating is configured to provide at least one thread means for processing inbound signals from the at least one gateway means being sent to the means for analyzing data and controlling actions via an inbound means and for processing outbound signals to the at least one gateway means received from the means for analyzing data and controlling actions via an outbound means, the inbound means and the outbound means operating asynchronously to provide non-blocking messaging.

DeKoning teaches a means for communicating is configured to provide at least one thread means for processing inbound signals from the at least one gateway means being sent to the means for analyzing data and controlling actions via an inbound means and for processing outbound signals to the at least one gateway means received from the means for analyzing data and controlling actions via an outbound means, the inbound means and the outbound means operating asynchronously to provide non-blocking messaging (see col.17, lines 8-16: "multi-threaded programming technique may be employed within a RAID controller to permit other processing to continue while element 700 is operable to await receipt of a new I/O request").

It would have been obvious to a person of ordinary skill in the art at the time the invention was made to modify the system of Guha in view of DeKoning by implementing

means for communicating is configured to provide at least one thread means for processing inbound signals from the at least one gateway means being sent to the means for analyzing data and controlling actions via an inbound means and for processing outbound signals to the at least one gateway means received from the means for analyzing data and controlling actions via an outbound means. One would be motivated to do so because Guha teaches of dynamically balancing input/output loads on the storage center (see pg.2, [0018]).

As per **claim 23**, Guha teaches a Service Level Agreement (SLA) server comprising:

means for managing and maintaining connections to the means for storing service level agreements (see pg.3, [0068]: "Based on the policy specified in the QoS Enforcer 34, content request may be dropped or requeued");

means for analyzing data and controlling actions based on the service level agreements and policy (see pg.4, [0045]: "The routing decisions are determined through a combination of preset QoS policy and current expected load at the data center");

means for communicating with clients to display monitoring information and for communication with the means for managing and maintaining connections to the means for storing service level agreement (see.Fig.5, #62); and

means for communicating with the at least one gateway means to collect data and send throttling requests based upon signals from means for analyzing data and

controlling actions (see pg.3, [0068]: "Based on the policy specified in the QoS Enforcer 34, content request may be dropped or requeued").

Guha does not explicitly teach wherein the means for communicating is configured to provide at least one thread pair for processing inbound signals from the at least one gateway means being sent to the means for analyzing data and controlling action via an inbound means and for processing outbound signals to the at least one gateway means received from the means for analyzing data and controlling action via an outbound means, the inbound means and the outbound means operating asynchronously to provide non-blocking messaging.

DeKoning teaches a means for communicating is configured to provide at least one thread pair for processing inbound signals from the at least one gateway means being sent to the means for analyzing data and controlling action via an inbound means and for processing outbound signals to the at least one gateway means received from the means for analyzing data and controlling action via an outbound means, the inbound means and the outbound means operating asynchronously to provide non-blocking messaging (see col.17, lines 8-16: "multi-threaded programming technique may be employed within a RAID controller to permit other processing to continue while element 700 is operable to await receipt of a new I/O request").

It would have been obvious to a person of ordinary skill in the art at the time the invention was made to modify the system of Guha in view of DeKoning so that means for communicating is configured to provide at least one thread pair for processing inbound signals from the at least one gateway means being sent to the means for

analyzing data and controlling action via an inbound means and for processing outbound signals to the at least one gateway means received from the means for analyzing data and controlling action via an outbound means. One would be motivated to do so because Guha teaches of dynamically balancing input/output loads on the storage center (see pg.2, [0018]).

As per **claim 24**, Guha teaches a performance monitor for controlling communication between two functional entities, comprising:

at least one processing means configured for processing signals between at least one I/O performance gateway and SLA services module processes (pg.5, [0069]: "communicated by the QoS Enforcer 34 to the content controller 36").

Guha does not explicitly teach a thread means associated with each of the at least one processors, each thread means processing inbound signals from the at least one I/O device/module being sent to the services module via an inbound means and for processing outbound signals to the at least one I/O device/module received from the services module via an outbound means, the inbound means and the outbound means operating asynchronously to provide non-blocking messaging.

DeKoning teaches a thread means associated with each of the at least one processors, each thread means processing inbound signals from the at least one I/O device/module being sent to the services module via an inbound means and for processing outbound signals to the at least one I/O device/module received from the services module via an outbound means, the inbound means and the outbound means

operating asynchronously to provide non-blocking messaging (see col.17, lines 8-16: "multi-threaded programming technique may be employed within a RAID controller to permit other processing to continue while element 700 is operable to await receipt of a new I/O request").

It would have been obvious to a person of ordinary skill in the art at the time the invention was made to modify the system of Guha in view of DeKoning by implementing a thread means associated with each of the at least one processors, each thread means processing inbound signals from the at least one I/O device/module being sent to the services module via an inbound means and for processing outbound signals to the at least one I/O device/module received from the services module via an outbound means. One would be motivated to do so because Guha teaches of dynamically balancing input/output loads on the storage center (see pg.2, [0018]).

DEPENDENT:

As per **claims 2 and 9**, which respectively depend on claims 1 and 8, Guha further teaches wherein the SLA services module further comprises:

a performance analyzer for setting throttling parameters and discovering new I/O performance gateways (see pg.5, [0069]);

an entity service module for providing in-memory caching of collected statistical data by polling data from the I/O performance gateways (see pg.3, [0038]); and

a policy manager for ensuring actions meet service level agreements and policy rules (see pg.4, [0045]).

As per **claim 3**, which depends on claim 1, Guha teaches of further comprising a storage resource manager for monitoring the storage system (see pg.4, [0050]).

As per **claims 5, 11, and 15**, which respectively depend on claims 1, 8, and 14, Guha and DeKoning further teaches wherein the performance monitor includes a plurality of processors for controlling communication (see Fig.1, #1, #2, #3), the at least one thread pair comprising a thread pair for each of the plurality of processors (DeKoning: see claim 1 rejection above).

As per **claims 7, 13, 17, and 20**, which respectively depend on claims 1, 8, 14 and 18, Guha and DeKoning further teaches wherein the providing at least one processor further comprises providing a single processor (see Fig.4), and wherein the providing a thread pair associated with each of the at least one processors further comprises providing a signal thread pair, wherein an inbound thread receives signals from all gateways and provides a path to all processes of the SLA services module (DeKoning: see claim 1 rejection above).

Allowable Subject Matter

6. Claims 4, 6, 10, 12, 16, and 19 are objected to as being dependent upon a rejected base claim, but would be allowable if rewritten in independent form including all of the limitations of the base claim and any intervening claims.

The following is an examiner's statement of reasons for allowance:

The prior art of record does not disclose, teach, or suggest neither singly nor in combination the claimed limitation of "wherein the SLA server further includes a mailbox

disposed between the SLA services module, the database manager, the application server and the performance monitor, the mailbox providing a non-blocking two-step communication scheme for allowing concurrent servicing of multiple I/O requests and database requests" as recited in **claims 4 and 10**.

The prior art of record does not disclose, teach, or suggest neither singly nor in combination the claimed limitation of "wherein the providing a thread pair associated with each of the at least one processors further comprises associating each thread pair associated with a processor with a subgroup of SLA processes of the SLA services module and a subgroup of gateways" as recited in **claims 6, 12, 16, and 19**.

Conclusion

7. For the reasons above, claims 1-3, 5, 7-9, 11, 13-15, 17, 18 and 20-24 have been rejected and claims 4, 6, 10, 12, 16, and19 have been objected. Claims 1-24 remain pending.

8. Any inquiry concerning this communication or earlier communications from the examiner should be directed to Michael Y. Won whose telephone number is 571-272-3993. The examiner can normally be reached on M-Th: 7AM-5PM.

If attempts to reach the examiner by telephone are unsuccessful, the examiner's supervisor, Saleh Najjar can be reached on 571-272-4006. The fax phone number for the organization where this application or proceeding is assigned is 571-273-8300.

Information regarding the status of an application may be obtained from the Patent Application Information Retrieval (PAIR) system. Status information for published applications may be obtained from either Private PAIR or Public PAIR. Status information for unpublished applications is available through Private PAIR only. For more information about the PAIR system, see http://pair-direct.uspto.gov. Should you have questions on access to the Private PAIR system, contact the Electronic Business Center (EBC) at 866-217-9197 (toll-free). If you would like assistance from a USPTO Customer Service Representative or access to the automated information system, call 800-786-9199 (IN USA OR CANADA) or 571-272-1000.

Michael Won

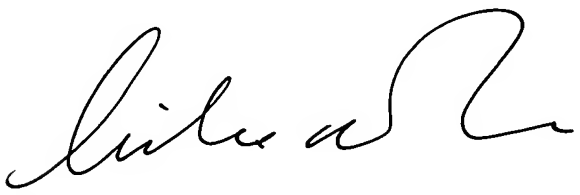

February 27, 2007